

UNITED STATES DISTRICT COURT
DISTRICT OF MASSACHUSETTS

CIVIL ACTION NO. 16-11613-RGS

EGENERA, INC.

v.

CISCO SYSTEMS, INC.

MEMORANDUM AND ORDER ON
DEFENDANT'S MOTION TO DISMISS

February 14, 2017

STEARNS, D.J.

The desire to economize time and mental effort in arithmetical computations, and to eliminate human liability to error, is probably as old as the science of arithmetic itself. This desire has led to the design and construction of a variety of aids to calculation, beginning with groups of small objects, such as pebbles, first used loosely, later as counters on ruled boards, and later still as beads mounted on wires fixed in a frame, as in the abacus.

— Howard Aiken, father of the Mark I IBM computer¹

Beginning with the invention by Blaise Pascal of the mechanical calculator, and culminating in our times with the integrated circuit-based computer, the ability of modern computers to aid human beings in performing tasks requiring the processing of large amounts of data has, as

¹ In Zenon W. Pylyshyn & Liam J. Bannon, *Perspectives on the Computer Revolution* (1989).

Gordon Moore predicted, grown exponentially as transistors have miniaturized, while doubling in capacity roughly every eighteen months since 1965. In 1874, Frank Stephen Baldwin was granted the first American patent (No. 153,522) for a calculating machine, the arithmometer. The number of “calculator patents” granted since is impossible to estimate accurately, but certainly runs to the hundreds of thousands. Not all of these patents are valid. Patents only protect inventions. They do not (or at least they are not supposed to) vest intellectual property rights in patentees who stake claims to the “building blocks of human ingenuity.” *Alice Corp. v. CLS Bank Int'l*, 134 S. Ct. 2347, 2354 (2014) (internal quotation marks and citation omitted).

One of the challenges for the patent system, as laid out by the Supreme Court in the recent *Alice* decision, is to separate out new and useful applications of abstract ideas from impermissible attempts to monopolize them.

Stating an abstract idea “while adding the words ‘apply it’” is not enough for patent eligibility. *Mayo Collaborative Servs. v. Prometheus Labs., Inc.*, 132 S. Ct. 1289, 1294 (2012). Nor is limiting the use of an abstract idea “to a particular technological environment.” *Bilski v. Kappos*, 561 U.S. 593, 610-611 (2010). Stating an abstract idea while adding the words “apply it with a computer” simply combines those two steps, with the same deficient result.

Alice, 134 S. Ct. at 2358.

In this intellectual property dispute, plaintiff Egenera, Inc., alleges that defendant Cisco Systems, Inc., infringes three of Egenera's patents – United States Patent No. 6,971,044 (the '044 patent), United States Patent No. 7,1178,059 (the '059 patent), and United States Patent No. 7,231,430 (the '430 patent), all of which set out claims to an improved enterprise computing system. Cisco moves to dismiss the Complaint, contending that Egenera's patents do not assert viable claims to patentable subject matter as required by 35 U.S.C. § 101. The court heard oral argument on the motion to dismiss on February 8, 2017.

The Asserted Patents

The '044 and the '430 patents arise from two patent applications filed on January 4, 2002. Both patents claim priority to the same provisional application filed on April 20, 2001, and share virtually identical specifications. The '044 patent is entitled “Service Cluster and Method in a Processing System with Failover Capability,” and was issued on November 29, 2005. It lists as its inventors Scott Geng, Pete Manca, Paul Curtis, Ewan Milne, Max Smith, Alan Greenspan, Edward Duffy, and Peter Schulter. The '430 patent is entitled “Reconfigurable, Virtual Processing System Cluster, Network, and Method,” and was issued on June 12, 2007. The '430 patent

lists the same inventors as the '044 patent, with three additions – Vern Brownell, Ben Sprachman, and Dan Busby.

The '044 and '430 patents are directed to solving problems in configuring, deploying, and maintaining enterprise and application servers.

For example, when deploying 24 conventional servers, more than 100 discrete connections may be required to configure the overall system. Managing these cables is an ongoing challenge, and each represents a failure point. Attempting to mitigate the risk of failure by adding redundancy can double the cabling, exacerbating the problem while increasing complexity and costs.

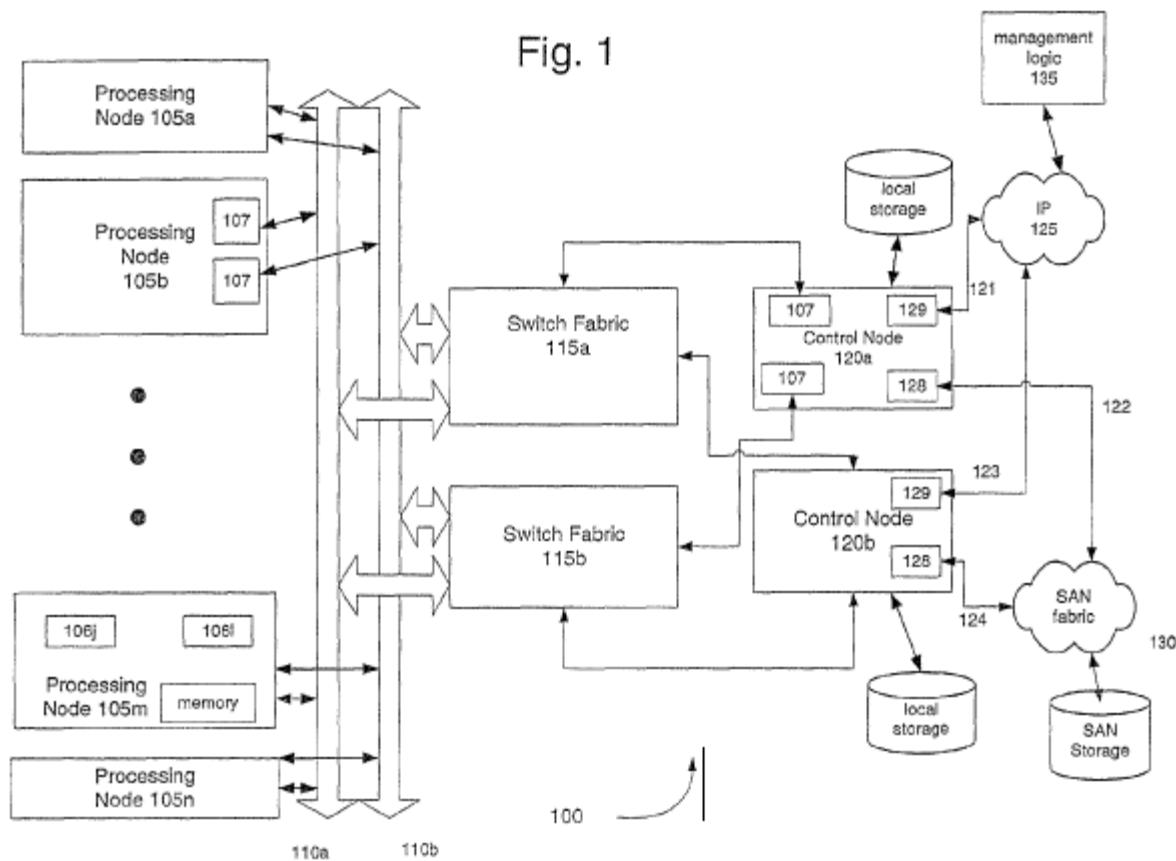
'044 patent, col. 1, ll. 41-47. Given that “personnel from multiple information technology (IT) functions (electrical, networking, etc.) must participate to deploy processing and networking resources . . . it can take weeks or months to deploy a new computer server.” *Id.* col. 1, ll. 22-27. In addition, to ensure high availability of server resources, “a failover server must be deployed for every primary server . . . [requiring] complex management software and professional services.” *Id.* col. 1, ll. 49-51. Finally, a post-deployment system adjustment “often requires a ‘forklift upgrade,’ meaning more hardware/software systems are added, needing new connections and the like.” *Id.* col. 1, ll. 56-58.

The '044 and '430 patents seek to improve on these limitations by creating a “processing platform from which virtual systems may be deployed through configuration commands.” *Id.* col. 2, ll.

The platform provides a large pool of processors from which a subset may be selected and configured through software commands to form a virtualized network of computers (“processing area network” or “processor clusters”) that may be deployed to serve a given set of applications or customer. The virtualized processing area network (PAN) may then be used to execute customer specific applications, such as Web-based server applications. The virtualization may include virtualization of local area networks (LANs) or the virtualization of I/O storage. By providing such a platform, processing resources may be deployed rapidly and easily through software via configuration commands, e.g., from an administrator, rather than through physically providing servers, cabling network and storage connections, providing power to each server and so forth.

Id. col. 2, l. 59 - col. 3, l. 7. Figure 1 illustrates the physical topography of a prototypical platform.

Fig. 1



As shown in FIG. 1, a preferred hardware platform 100 includes a set of processing nodes 105a-n connected to a switch fabrics 115a,b via high-speed, interconnect 110a,b. The switch fabric 115a,b is also connected to at least one control node 120a,b that is in communication with an external IP network 125 (or other data communication network), and with a storage area network (SAN) 130. A management application 135, for example, executing remotely, may access one or more of the control nodes via the IP network 125 to assist in configuring the platform 100 and deploying virtualized PANs.

Under certain embodiments, about 24 processing nodes 105a-n, two control nodes 120, and two switch fabrics 115a,b are contained in a single chassis and interconnected with a fixed, pre-wired mesh of point-to-point (PtP) links. Each processing node 105 is a board that includes one or more (e.g., 4) processors 106j-l, one or more network interface cards (NICs) 107, and local memory (e.g., greater than 4 Gbytes) that, among other things, includes some BIOS firmware for booting and initialization. There is no local disk for the processors 106; instead all storage, including storage needed for paging, is handled by SAN storage devices 130.

Each control node 120 is a single board that includes one or more (e.g., 4) processors, local memory, and local disk storage for holding independent copies of the boot image and initial file system that is used to boot operating system software for the processing nodes 105 and for the control nodes 106. Each control node communicates with SAN 130 via 100 megabyte/second fibre channel adapter cards 128 connected to fibre channel links 122, 124 and communicates with the Internet (or any other external network) 125 via an external network interface 129 having one or more Gigabit Ethernet NICs connected to Gigabit Ethernet links 121,123. (Many other techniques and hardware may be used for SAN and external network connectivity.) Each control node includes a low speed Ethernet port (not shown) as a dedicated management port, which may be used instead of remote, web-based management via management application 135.

The switch fabrics is composed of one or more 30-port Giganet switches 115, such as the NIC-CLAN 1000 and clan 5300 switch, and the various processing and control nodes use corresponding NICs for communication with such a fabric module. Giganet switch fabrics have the semantics of a Non-Broadcast Multiple Access (NBMA) network. All inter-node communication is via a switch fabric. Each link is formed as a serial connection between a NIC 107 and a port in the switch fabric 115. Each link operates at 112 megabytes/second.

'044 patent, col. 3, ll. 9-59.

The platform is designed to utilize virtual MAC addresses² to respond to failovers.

A plurality of computer processors are [sic] connected to an internal communication network. A virtual local area communication network over the internal network is defined and established. Each computer processor in the virtual local area communication network has a corresponding virtual MAC address and the virtual local area network provides communication among a set of computer processors but excludes the processors from the plurality not in the defined set. A virtual storage space is defined and established with a defined correspondence to the address space of the storage network. In response to a failure by a computer processor, a computer processor from the plurality is allocated to replace the failed processor. The MAC address of the failed processor is assigned to the processor that replaces the failed processor. The virtual storage space and defined correspondence of the failed processor is assigned to the processor that replaces the failed processor. The virtual local area network is reestablished to include the processor that replaces the failed processor and to exclude the failed processor.

² MAC (media access control) addresses are unique identifiers permanently assigned to network interface hardware components such as Ethernet cards and routers.

The '430 patent lists 8 claims, of which claim 1 is asserted.

1. A platform for automatically deploying at least one virtual processing area network, in response to software commands, said platform comprising:

a plurality of computer processors connected to an internal communication network;

at least one control node in communication with an external communication network and in communication with an external storage network having an external storage address space, wherein the at least one control node is connected to the internal communication network and thereby in communication with the plurality of computer processors, said at least one control node including logic to receive messages from the plurality of computer processors, wherein said received messages are addressed to the external communication network and to the external storage network and said at least one control node including logic to modify said received messages to transmit said modified messages to the external communication network and to the external storage network;

configuration logic for receiving and responding to said software commands, said software commands specifying (i) a number of processors for a virtual processing area network (ii) a virtual local area network topology defining interconnectivity and switching functionality among the specified processors of the virtual processing area network, and (iii) a virtual storage space for the virtual processing area network, said configuration logic including logic to select, under programmatic control, a corresponding set of computer processors from the plurality of computer processors, to program said corresponding set of computer processors and the internal communication network to establish the specified virtual local area network topology, and to program the at least one control node to define a

virtual storage space for the virtual processing area network, said virtual storage space having a defined correspondence to a subset of the external storage address space of the external storage network; and

wherein the plurality of computer processors and the at least one control node include network emulation logic to emulate Ethernet functionality over the internal communication network.

The '044 patent lists 6 claims, of which claim 1 is asserted.

1. A platform for computer processing, connectable to an external communication network and a storage network and comprising:

a plurality of computer processors connected to an internal communication network;

configuration logic to define and establish (a) a virtual local area communication network over the internal network, wherein each computer processor in the virtual local area communication network has a corresponding virtual MAC address and the virtual local area network provides communication among a set of computer processors but excludes the processors from the plurality not in the defined set, and (b) a virtual storage space with a defined correspondence to the address space of the storage network; and

failover logic, responsive to a failure of a computer processor, to allocate a computer processor from the plurality to replace the failed processor, the failover logic including logic to assign the virtual MAC address of the failed processor to the processor that replaces the failed processor, logic to assign the virtual storage space and defined correspondence of the failed processor to the processor that replaces the failed processor, and logic to reestablish the virtual local area network to include the processor that replaces the failed processor and to exclude the failed processor.

The '059 patent derives from an application filed on May 7, 2003. It is entitled "Disaster Recovery for Processing Resources Using Configurable Deployment Platform," and was issued on February 13, 2007. It lists as inventors Alan Greenspan, Borne Goodman-Mace, Michael Johnson, Siping Liu, and Claude Keswani. The '059 patent discloses a disaster recovery system that can be configured to mimic the topology of a failed primary system. Most modern enterprise computing systems have intricate internal interconnections among computers and components that perform a host of differentiated tasks. In the typical system,

[a] firewall acts as an interface to the Internet, for example, to receive various requests therefrom. The firewall communicates with a load balancer, which attempts to distribute the processing load on the overall system among a multiplicity of processing nodes. For example, the load balancer may distribute requests among a multiplicity of Web servers. Each Web server, in turn, may perform some analysis of a task it receives and invoke an appropriate application server. Each application server may in turn interact with database or file server. Each of the various entities may be executing on its own respective processing or server node. . . . Adding to the complication . . . are the various hubs, switches, cabling, and the like necessary to create the depicted processing network. Moreover, various versions of software may be executing.

'059 patent, col. 2, ll. 36-55.

To date, processor-side aspects of disaster recovery have largely been handled by requiring processing resources on the secondary site to be identical to those of the first site and to wait in standby mode. This is complicated and costly, as suggested by

the complexity of the multi-tiered architecture. Moreover, modern processor networks are often changed for a variety of reasons. If such a network is a primary site network, then the changes also need to be made to the secondary, or else the enterprise risks that its disaster recovery system will not work as expected.

Id. col. 1, l. 61 - col. 2, l. 3.

The '059 patent eliminates the need for a standby secondary system that is physically identical to the primary system by creating a configurable PAN platform.

Processing resources and configuration at the primary site are characterized into a specification with a defined set of variables, and the specification is stored in a secure way. The set of information that characterizes the resources (i.e., the resource's "personality") includes information such as the number of processing area networks (PANs) at the primary site, for each such PAN the number of nodes that should be allocated, the network connectivity among processors, storage mappings and the like The failover site uses a software configurable platform that allows one or more independent processing networks to be deployed (or instantiated) in response to software commands. . . . The configuration specification is accessed and used to issue a set of commands on the configurable platform to instantiate processing resources on the failover site consistent with the specification.

Id. col. 3, ll. 24-42. The '059 patent lists 10 claims, of which claim 10 is asserted:

10. A system of providing processing resources to respond to a fail-over condition in which a primary site includes a configuration of processing resources, comprising:

a computer-readable specification that describes a configuration of processing resources of the primary site;

a configurable processing platform capable of deploying processing area networks in response to software commands;

logic to generate software commands to the configurable platform to deploy processing resources corresponding to the specification;

wherein the processing resources at the primary site include a plurality of independent processing area networks and wherein the specification describes all of the independent processing area networks.

ANALYSIS

35 U.S.C. § 101 provides that “[w]hoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.” Although the scope of patentable subject matter is “expansive,” laws of nature, physical phenomena, and abstract ideas have long been held to be patent-ineligible. *Diamond v. Chakrabarty*, 447 U.S. 303, 308-309 (1980). More recently, in addressing § 101 patentability, the Supreme Court has outlined a two-step query to be used in analyzing subject matter eligibility.

First, we determine whether the claims at issue are directed to one of those patent-ineligible concepts. [Mayo,] 132 S. Ct. at 1296-1297. If so, we then ask, “[w]hat else is there in the claims before us?” *Id.*, [] 132 S. Ct., at 1297. To answer that question,

we consider the elements of each claim both individually and “as an ordered combination” to determine whether the additional elements “transform the nature of the claim” into a patent-eligible application. *Id.*, [] 132 S. Ct., at 1298, 1297. We have described step two of this analysis as a search for an “inventive concept”— *i.e.*, an element or combination of elements that is “sufficient to ensure that the patent in practice amounts to significantly more than a patent upon the [ineligible concept] itself.” *Id.*, [] 132 S. Ct., at 1294.

Alice Corp., 134 S. Ct. at 2355.

Cisco argues that Egenera’s asserted claims fail both prongs of the *Alice* paradigm because they are directed to abstract ideas, and do not disclose any inventive concepts. The arguments are heavily intertwined. “Recent cases . . . suggest that there is considerable overlap between step one and step two [of *Alice*], and in some situations this analysis could be accomplished without going beyond step one.” *Amdocs (Israel) Ltd. v. Openet Tel., Inc.*, 841 F.3d 1288, 1294 (Fed. Cir. 2016). Although the Federal Circuit has broadly stated that “[t]he abstract idea exception prevents patenting a result where ‘it matters not by what process or machinery the result is accomplished,’” *McRO, Inc. v. Bandai Namco Games Am. Inc.*, 837 F.3d 1299, 1312 (Fed. Cir. 2016), quoting *O’Reilly v. Morse*, 56 U.S. (15 How.) 62, 113 (1854); there is “no [] single, succinct, usable definition or test” to determine exactly what an “abstract idea” encompasses. *Amdocs (Israel)*, 841 F.3d at 1294. “Instead of a definition, then, the decisional mechanism courts now apply is to examine

earlier cases in which a similar or parallel descriptive nature can be seen – what prior cases were about, and which way they were decided.” *Id.* Claims directed to “improvements in computer-related technology are [not] inherently abstract.” *Enfish, LLC v. Microsoft Corp.*, 822 F.3d 1327, 1335 (Fed. Cir. 2016). The most relevant inquiry is “whether the claims are directed to an improvement to computer functionality versus being directed to an abstract idea.” *Id.*

The '430 and '044 Patents

Cisco neatly sums up its argument by characterizing the '430 and '044 patents as being directed to “setting up a logical group of resources within a physically connected group of resources” – “the archetypal abstract idea” of “replacing manual functionality with software.” Cisco Br. at 17.

The [] claims purport to apply generic software “logic” (claimed in purely generic functional terms) to accomplish tasks through software that had long been accomplished through physical cabling. Although network engineers have long connected computers to each other and other network equipment with physical cabling to form a group within a network, the patents purport to use unspecified, functionally described “logic” to replace manual functionality.

Id. Cisco likens the '430 and '044 patents to others that have failed the patentability test because they offer no more than “an abstract idea implemented in a particular technological environment.” *Appistry, Inc. v. Amazon.com, Inc.*, 2015 WL 4210890, at *2 (W.D. Wash. July 9, 2015)

(*Appistry I*). In *Appistry*, the court found the disputed claims unpatentable because they

recite the abstract idea of distributed processing akin to the military’s command and control system, a longstanding and intuitive practice used by many large hierarchical organizations that value speed, efficiency, reliability, and accountability. The patents describe systems and methods of using a network of multiple actors to efficiently and reliably process information and/or complete a task by breaking down the job into small pieces, each handled by a different actor organized within an internal hierarchy.

Id.; see also *Appistry, Inc. v. Amazon.com Inc.*, 2016 WL 3906905, at *3 (W.D. Wash. July 19, 2016) (*Appistry II*) (concluding that “virtually identical claims” are also directed to the same unpatentable abstract concept).

Likewise, in *Intellectual Ventures I LLC v. Capital One Bank (USA)*, 792 F.3d 1363, 1367 (Fed. Cir. 2015), the Federal Circuit found claims directed to the computerized “tracking of financial transactions to determine whether they exceed a pre-set spending limit (i.e., budgeting)” to be unpatentably abstract. In another similar case, *In re TLI Communications LLC Patent Litigation*, 823 F.3d 607 (Fed. Cir. 2016), although the disputed claim referenced tangible computer components, the Court found it invalid where “the recited physical components merely provide a generic environment in which to carry out the abstract idea of classifying and storing digital images in an organized manner.” *Id.* at 611.

It is well-settled that mere recitation of concrete, tangible components is insufficient to confer patent eligibility to an otherwise abstract idea. Rather, the components must involve more than performance of “well-understood, routine, conventional activit[ies]” previously known to the industry.” *Alice*, 134 S. Ct. at 2359 (quoting *Mayo*, 132 S. Ct. at 1294).

Id. at 613.

Egenera, for its part, maintains that the claims of the ’430 and ’044 patents are directed to concrete enhancements of computer functionality.

[C]laim 1 [of the ’430 patent] recites a specific computing platform for deploying PANs – which are further described in the specification – that includes, among other elements, pools of processors, interfaces and/or connections to internal and external communication networks, control nodes with specific capabilities, and configuration logic for receiving and responding to software commands that specify “a virtual local area network topology defining interconnectivity and switching functionality among the specified processors of the virtual processing area network.”

Egenera Opp’n at 8. According to Egenera, the asserted claims are near relations of the patentable claims described in *Enfish* and *McRO*. In *Enfish*, the Federal Circuit rejected the district court’s conclusion that claims directed to self-referential data tables are unpatentably abstract. 822 F.3d at 1337-1339.

In sum, the self-referential table recited in the claims on appeal is a specific type of data structure designed to improve the way a computer stores and retrieves data in memory. The specification’s disparagement of conventional data structures, combined with language describing the “present invention” as including the features that make up a self-referential table,

confirm that our characterization of the “invention” for purposes of the § 101 analysis has not been deceived by the “draftsman’s art.” *Cf. Alice*, 134 S. Ct. at 2360. In other words, we are not faced with a situation where general-purpose computer components are added post-hoc to a fundamental economic practice or mathematical equation. Rather, the claims are directed to a specific implementation of a solution to a problem in the software arts.

Id. at 1339. In *McRO*, the Federal Circuit determined that claims setting out a specific set of rules to improve computer animation of facial expressions were not unpatentably abstract. 837 F.3d at 1313-1316.

[T]he automation goes beyond merely “organizing [existing] information into a new form” or carrying out a fundamental economic practice. *Digitech [Image Techs., LLC v. Elecs. For Imagine, Inc.]*, 758 F.3d [1344,] 1351 [(Fed. Cir. 2014)]; *see also Alice*, 134 S. Ct. at 2356. The claimed process uses a combined order of specific rules that renders information into a specific format that is then used and applied to create desired results: a sequence of synchronized, animated characters. While the result may not be tangible, there is nothing that requires a method “be tied to a machine or transform an article” to be patentable. *Bilski [v. Kappos]*, 561 U.S. [593,] 603 [(2010)] (discussing 35 U.S.C. § 100(b)). The concern underlying the exceptions to § 101 is not tangibility, but preemption. *Mayo*, 132 S. Ct. at 1301.

Id. at 1315.

It takes no great leap of imagination to grasp that the ’430 and ’044 patents, like the patents in *Enfish* and *McRO*, are directed to the solution of a real problem by the creation of an innovative processing platform with a positive result, namely an improved enterprise computing system. The Federal Circuit has cautioned judges “[to] be careful to avoid

oversimplifying the claims’ by looking at them generally and failing to account for the specific requirements of the claims,” *McRO*, 837 F.3d at 1313 (citation omitted), lest “the exceptions to § 101 swallow the rule.” *Enfish*, 822 F.3d at 1337. While a virtual computer network may be seen from twenty thousand feet as simply “a logical group of resources within a physically connected group of resources,” the ’430 and ’044 patents do not claim virtual networks and do not monopolize the underlying idea.³ Rather, the ’430 patent is directed to a computing platform that is capable of “automatically deploying at least one virtual processing area network” with a specified topology. *See ’430 patent, claim 1* (“configuration logic for receiving and responding to . . . a virtual area network topology defining interconnectivity and switching functionality”).

Claim 1 of the ’430 patent describes the physical composition of the platform. In addition to “a plurality of processors connected to an internal

³ Cisco’s other generalization – that the claimed platform replaces the conventional manual activity of cabling computers with unspecified software logic – is also inaccurate. As the claim language and the recited portion of the specification makes clear, the claimed system depends on extensive physical connections between its component processors. *See ’044 patent, col. 3, ll. 9-59; see also Amdocs (Israel)*, 841 F.3d at 1299 (§ 101 subject matter eligibility to be considered in light of the specification). Claim 2 of the ’430 patent, specifying a point-to-point switch fabric as the internal communication network, is further evidence that the specific configuration of hardware components is essential to the functionality of the platform.

communication network,” the platform has “at least one control node in communication with an external communication network[,] . . . an external storage address space, [and] the internal communication network.” Unlike the claims found unpatentable in *Appistry*, *Intellectual Ventures*, and *TLI*, the recitation of hardware components and their connections is not “a situation where general-purpose computer components are added post-hoc to a fundamental economic practice or mathematical equation.” *Enfish*, 822 F.3d at 1339. Claim 1 requires the components of the system to perform certain specialized functions to accomplish the stated goal. The “at least one control node” must facilitate modified messaging between the processors and the external networks. *See* ’430 patent, claim 1 (“said at least one control node including logic to receive messages from the plurality of computer processors . . . addressed to the external communication [and] storage network[s] . . . [and] logic to modify said received messages to transmit . . . to the external . . . network[s]”). The platform must map physical processors to emulate the members of a virtual network with a particular topology, and map physical external storage to virtual storage. *See id.* (“logic to select . . . a corresponding set of computer processors . . . , to program said corresponding set of computer processors and the internal communication network to establish the specified virtual local area network topology, and to

program at least one control node to define a virtual storage space for the virtual processing area network”). Finally, the processors and the control node(s) must be capable of “emulat[ing] Ethernet functionality over the internal communication network.” *Id.*

The ’044 patent is directed to a computer platform that can substitute for a failed processor of a virtual network without having to physically replace the processor itself. Claim 1 of the ’044 patent requires a plurality of connected processors to establish a virtual network with a virtual storage space; and, by using virtual MAC addresses, to replace a failed processor by remapping the assigned address and storage. *See ’044 patent, claim 1 (“logic to assign the virtual MAC address of the failed processor to the processor that replaces the failed processor, logic to assign the virtual storage space and defined correspondence of the failed processor to the processor that replaces the failed processor”).*

That certain aspects of the functionality of the claimed platforms is specified through programming logic is not an impediment to subject matter eligibility. In *McRO*, in approving claims directed to a software method for the improved computer animation of facial expressions, the Federal Circuit noted that there was no tangibility requirement for patent eligibility. 837 F.3d at 1315. Nor is it a bar if the individual elements of the claimed systems

are, as Cisco contends, routine and conventional in the art. In *Bascom Global Internet Servs., Inc. v. AT&T Mobility LLC*, 827 F.3d 1341 (Fed. Cir. 2016), the Federal Circuit found claims for an internet content filtering system patent-eligible despite being directed to the abstract concept of content filtering and being composed of generic computer components and conventional methods. *Id.* at 1349-1350.

The inventive concept inquiry requires more than recognizing that each claim element, by itself, was known in the art. As is the case here, an inventive concept can be found in the non-conventional and non-generic arrangement of known, conventional pieces.

The inventive concept described and claimed in the '606 patent is the installation of a filtering tool at a specific location, remote from the end-users, with customizable filtering features specific to each end user. This design gives the filtering tool both the benefits of a filter on a local computer and the benefits of a filter on the ISP server. BASCOM explains that the inventive concept rests on taking advantage of the ability of at least some ISPs to identify individual accounts that communicate with the ISP server, and to associate a request for Internet content with a specific individual account. . . . According to BASCOM, the inventive concept harnesses this technical feature of network technology in a filtering system by associating individual accounts with their own filtering scheme and elements while locating the filtering system on an ISP server. *See Research Corp. Techs. v. Microsoft Corp.*, 627 F.3d 859, 869 (Fed. Cir. 2010) (“[I]nventions with specific applications or improvements to technologies in the marketplace are not likely to be so abstract that they override the statutory language and framework of the Patent Act.”). On this limited record, this specific method of filtering Internet content cannot be said, as a matter of law, to have been conventional or generic.

Id. at 1350.

Like the claims in *Bascom*, claim 1 of the '430 patent offers a useful, non-conventional, and non-generic arrangement of components and functionality, namely, a unique arrangement of networked processors and control nodes with the specified messaging, mapping, and Ethernet emulation functions, *see id.*, that is “directed to a specific implementation of a solution to a problem.” *Enfish*, 822 F.3d at 1339; *see also McRO*, 837 F.3d at 1515. At this early stage of the litigation, Cisco has not shown *at the time of invention* that the combination of hardware and functionality taught by the '430 and '044 patents was either conventional or generic.⁴ *See Bascom*, 827 F.3d at 1350.

⁴ Although there is some overlap between the § 101 eligibility analysis and an anticipation or obviousness analysis under 35 U.S.C. §§ 102 & 103 in determining whether claim elements are “routine and conventional,”

the concept of inventiveness is distinct from that of novelty. Novelty is the question of whether the claimed invention is new. Inventiveness is the question of whether the claimed matter is invention at all, new or otherwise. The inventiveness inquiry of § 101 should therefore not be confused with the separate novelty inquiry of § 102 or the obviousness inquiry of § 103. Accordingly, the Supreme Court has cautioned that “[t]he obligation to determine what type of discovery is sought to be patented must precede the determination of whether that discovery is, in fact, new or obvious.”

Amdocs (Israel), 841 F.3d at 1311. To the extent that the subject matter eligibility inquiry requires assessing the state of the art, that assessment, as in the §§ 102 and 103 context, must be conducted from the perspective of a

Like the self-referential data table of *Enfish* and the animation rules of *McRO*, the claimed processing platform presents an improvement in computer functionality. In addition to expediting system deployment, the platform removes a system's dependence on *specific* physical connections between processors while maintaining the desired performance. The ability to automatically deploy a virtual processing area network also provides efficiency, flexibility, and scalability not available in a manually cabled system. *See, e.g.*, '430 patent, col. 3, ll. 53-54 (“Under software control, the platform supports multiple, simultaneous and independent processing area networks.”); col. 28, ll. 50-55 (“It will be appreciated that deployment may be based on programmatic control. For example, more processors may be deployed under software control during peak hours of operation for the PAN, or corresponding more or less storage space for a PAN may be deployed under software algorithmic control.”).

Similarly, the '044 patent discloses a network of processors that establishes a virtual network utilizing reassignable virtual MAC addresses to support the event of a node failure. Although Cisco argues that virtual MAC addresses do no more than conventional MAC addresses in identifying a

person of ordinary skill in the art at the time of the invention to avoid the usual temptations of hindsight.

specific node on a network, *see Cisco Reply at 15*, there is a crucial difference. Conventional MAC addresses are unique and fixed to a specific component of network hardware. In a physical world, if the system fails, in addition to replacing the failed processor, the address resolution protocol (ARP⁵) tables have to be updated to reflect the new MAC address of the replacement hardware. The '044 patent removes the necessity of other computers relearning the new MAC address.

When in filtered mode, there will be one externally visible MAC address to which external nodes transmit packets for a set of virtual network interfaces. If that adapter goes down, then not only do the virtual network interfaces have to fail over to the other control node, but the MAC address must fail over too so that external nodes can continue to send packets to the MAC address already in the ARP caches. Under one embodiment of the invention, when a failed control node recovers, a single MAC address is manipulated and the MAC address does not have to be remapped on recovery.

'044 patent, col. 19, l. 61 - col. 20, l. 4. Whether at *Alice* step 1 or step 2, because the '430 and '044 patents are directed to systems that improve computer functionality, they claim patent-eligible subject matter.

⁵ The ARP is a network protocol that maps network (IP) addresses to a physical (MAC) hardware address to ensure that messages are delivered to the appropriate and intended recipient computer.

The '059 Patent

Cisco contends that the '059 patent is directed to the abstract idea of “using a ‘disaster plan’ to set up a backup site,” and discloses nothing more than “a generic primary site, a generic configurable platform located at a backup site, and a generic specification made available (via storage system) to the backup site.” Cisco Br. at 7. According to Cisco, claim 10 of the '059 patent is “highly analogous” to the claim found ineligible in *Tranxition v. Lenovo (United States) Inc.*, 2016 WL 6775967 (Fed. Cir. Nov. 16, 2016). In *Tranxition*, the disputed claim “instruct[ed] a practitioner to (1) provide configuration information, (2) generate an extraction plan, (3) extract the configuration settings, (4) generate a transition plan, and (5) transition those settings to a new computer.” *Id.*, at *3. The Federal Circuit held the claim to be directed to the abstract idea of data migration.

Put another way, the stated aim of the patent is to automate the migration of data between two computers. This is not sufficient under step one of *Alice*. Contrary to *Tranxition*’s argument, the claim is not directed to an improvement to computer functionality. There is nothing in the claim to suggest that, once settings have been transitioned, the target computer will be any more efficient.

Id., at *3. Further, although a computer was able to accomplish the task more quickly and thoroughly, the claimed method applied the abstract idea

no differently than one would apply it manually. Consequently, the Court held that the claim did not disclose an inventive concept.

[I]t is not relevant that a human may perform a task differently from a computer. It is necessarily true that a human might apply an abstract idea in a different manner from a computer. What matters is the application. “Stating an abstract idea while adding the words ‘apply it with a computer’” will not render an abstract idea non-abstract. *See [Alice, 134 S. Ct.] at 2359.* There must be more.

Id., at *3.

In response, Egenera maintains that the '059 patent claims eligible subject matter because the '059 patent discloses a system for launching a secondary backup site on “a configurable processing platform capable of deploying processing area networks in response to software commands” corresponding to a “specification describ[ing] all of the independent processing area networks” of the primary site. *See '059 patent, claim 10.* Emphasizing the '059 patent’s extensive discussions of the platform and the specifications, Egenera argues that, like the self-referential data table of *Enfish* and animation rules of *McRO*, the '059 patent claims offer an improvement in computer functionality.

I do not think so. The claims of the '059 patent are set out at a level of generality that is indistinguishable from the underlying abstract idea. In *Internet Patents Corp. v. Active Network, Inc.*, 790 F.3d 1343, 1348 (Fed.

Cir. 2015), the Federal Circuit invalidated a claim directed to retaining user-entered information that is often lost in navigating online forms because the critical “maintaining state” limitation is “dissociated from any method by which maintaining the state is accomplished.” Similarly, the key functionality of claim 10 of the ’059 patent – “logic to generate software commands to the configurable platform to deploy processing resources corresponding to the specification” – is claimed in functional terms that restate the goal of the invention and “contains no restriction on how the result is accomplished.” *Id.*

Moreover, the inclusion of the processing area network limitations does not, as Egenera argues, serve to impart the necessary inventive concept. While the court agrees with Egenera that the processing area network platform of the ’430 and ’044 patents is patentable subject matter, unlike the claims of those two patents, the claim limitations of the ’059 patent do not correspond to any physical or functional aspect of the described system. The coupling of a generic step of setting up a disaster recovery backup site with a system capable of deploying processing area networks does no more than implement “an abstract idea [] in a particular technological environment.” *Appistry (I)*, 2015 WL 4210890, at *2. This does not satisfy § 101.

ORDER

For the forgoing reasons, Cisco's motion to dismiss is DENIED IN PART with respect to the '430 and '044 patents, and ALLOWED IN PART with respect to the '059 patent. The parties are requested to submit a joint proposed pre-trial schedule, consistent with L.R. 16.6, no later than Feb. 28, 2017.

SO ORDERED.

/s/ Richard G. Stearns

UNITED STATES DISTRICT JUDGE